

CLAIMS

What is claimed is:

1. A method for reorganizing a database table online, the method steps comprising:
 - a) moving a subset of records within the database table;
 - b) flagging each moved record as a reorganization record;
 - c) creating a reorganization pointer record for each moved record at the initial location of the moved record, the reorganization pointer record pointing to the new location of the moved record; and
 - d) establishing scanner process constraints based on whether a scanner process is commenced prior to or after the moving.
2. The method of claim 1 wherein the establishing further comprises:
 - d1) identifying table and index scanner processes commencing prior to the completion of the moving as old table scanners wherein old table scanners are constrained to access records in the database table by following reorganization pointers to access reorganization records and wherein old table scanners are constrained to ignore reorganization records when accessing records sequentially in a scan of the database table; and
 - d2) identifying table and index scanner processes commencing after the completion of the moving as new table scanners, wherein new table scanners are constrained to ignore reorganization pointers in accessing records in the database table.

1 3. The method of claim 2 further comprising the steps of:

2 e) deleting the reorganization pointer records and removing the
3 reorganization record flag from the records flagged after each query associated with an
4 old table scanner is complete; and

5 f) repeating steps a-d after the completion of each query associated with a
6 new scanner until the records in the database table meet a reorganization criteria.

1 4. The method of claim 3 in which step f) alternates between moving records from a
2 one of successively defined portions of the table and moving records in a predefined
3 order into a one of the successively defined vacated portions of the table.

1 5. The method of claim 4 in which the predefined order for moving records is
2 defined by a clustering index associated with the table.

1 6. The method of claim 4 in which the table further comprises overflow pointer
2 records pointing to overflow records and in which the method steps further comprise
3 steps defining each overflow pointer record in a one of the successively defined portions
4 of the table to be a reorganization overflow pointer record, and flagging the associated
5 overflow record to be a reorganization record.

1 7. The method of claim 4 in which the method steps further comprise the step of
2 preventing the allocation of space for new or updated records in each respective one of
3 the successively defined portions of the table when moving records from the one of

successively defined portions of the table.

8. The method of claim 3 wherein the deleting further comprises requesting and waiting for the exclusive availability of a latch which is shareable by each old table scanner process and which each old table scanner process obtains before commencing scanning and releases after the query associated with the old table scanner process has completed processing.

9. The method of claim 8 in which the latch is shareable by each index scanner process and each index scanner process obtains the latch before commencing scanning and releases the latch after the query associated with the index scanner process has completed processing.

10. The method of claim 3 in which the method steps further comprise steps to maintain the identity and status of moved records in a record identifier mapping table.

11. The method of claim 10 in which the record identifier mapping table comprises a mapping record corresponding to each of the moved records, each mapping record comprising an old record identifier for the initial position of the moved record, a new record identifier for the moved position of the moved record, and a status flag for the moved record.

12. The method of claim 11 in which the status flag for the moved record comprises

2 flags to indicate the completion of the move of the record, the completion of the clean up
3 for the record, and whether the record has been deleted.

1 13. The method of claim 10 in which the record identifier mapping table comprises a
2 hash table implemented having a hash on the new record identifier, and a hash table
3 containing a reverse record identifier mapping.

1 14. The method of claim 3 in which the method steps further comprise logging steps
2 for retaining data values relating to:

- 3 1) the start of step a),
4 2) the move of a normal record,
5 3) the move of a pointer record where, when moving records from a one of
6 successively defined portions of the table, the associated overflow record is not in the
7 vacate range,
8 4) the move of a pointer record where, when moving records from a one of
9 successively defined portions of the table, the associated overflow record is in the vacate
10 range,
11 5) the move of an overflow record during a vacate phase,
12 6) the successful move of a record,
13 7) the successful clean up of a row that has been moved, and
14 8) the completion of the online reorganization,
15 wherein the table may be rolled back without the loss of data due to the
16 reorganization of the table.

1 15. The method of claim 3 in which the method steps further comprise the step of
2 imposing a key value lock where an index scanner is accessing an ordered list of key
3 values in an index associated with the table, the key value lock preventing index values
4 from being updated while the index scanner accesses the ordered list of key values.

1 16. The method of claim 3 in which the method steps further comprise the step of
2 synchronising between scanners by imposing a share lock on records subject to the
3 moving and records subject to the deleting.

1 17. The method of claim 3 in which the method steps further comprise an idle step in
2 which neither the moving nor the deleting is performed.

1 18. A method for reorganizing a database table the method comprising the steps of:
2 a) a vacate move step;
3 b) a vacate clean up step;
4 c) a fill move step; and
5 d) a fill clean up step wherein each clean up step is synchronised to
6 commence at the completion of a move step and to commence only when all queries
7 launching scanner processes commenced before the completion of a move step have
8 completed and each move step is synchronised to commence at the completion of a clean
9 up step and to commence only when all queries launching scanner processes after the
10 completion of a previous move step have completed.

1 19. The method of claim 18 wherein step a) comprises:

- 2 i) moving each data record in a defined set of pages to a new location in a
3 page in the database table outside the defined set of pages;
- 4 ii) flagging each moved non-overflow data record as a reorganization record;
- 5 iii) creating a reorganization pointer record for each moved non-overflow data
6 record at the initial location of the moved record, the reorganization pointer record
7 pointing to the new location of the moved record;
- 8 iv) for each overflow pointer record located in the defined set of pages and
9 associated with a moved overflow data record, creating a reorganization pointer record
10 for the moved overflow data record at the location of the said reorganization pointer
11 record and flagging the moved overflow record as a reorganization record; and
- 12 v) flagging each overflow pointer record pointing to an overflow record
13 outside the defined set of pages to be a reorganization pointer record, and flagging each
14 of said overflow records as a reorganization record.

1 20. The method of claim 19 wherein step b) comprises:

- 2 i) deleting reorganization pointers in the defined set of pages; and
- 3 ii) removing the reorganization record flag from each reorganization record
4 in the database table.

1 21. The method of claim 20 wherein step c) comprises:

- 2 i) moving selected records to the defined set of pages such that the relative
3 ordering of the selected records matches the predefined ordering constraint;

- 4 ii) flagging each moved record as a reorganization record;
- 5 iii) creating a reorganization pointer record for each moved non-overflow
- 6 record at the initial location of the moved record, the reorganization pointer record
- 7 pointing to the new location of the moved record;
- 8 iv) creating a reorganization pointer record for each moved overflow record at
- 9 the location of the overflow pointer record for the moved record, the reorganization
- 10 pointer record pointing to the new location of the moved record.

1 22. The method of claim 21 wherein step d) comprises:

- 2 i) deleting reorganization pointers in the defined set of pages; and
- 3 ii) removing the reorganization record flag from each reorganization record
- 4 in the database table.

1 23. The method of claim 22, wherein table scanner processes commencing during

2 each move step are identified as old table scanners, wherein old table scanners are

3 constrained to access records in the database table by following reorganization pointers to

4 access reorganization records and are further constrained to access reorganization records

5 sequentially by ignoring reorganization records and table scanner processes commencing

6 after the completion of a move step are identified as new table scanners, wherein new

7 table scanners ignore reorganization pointers in accessing records in the database table.

1 24. The method of claim 22 in which each clean up synchronisation step comprises

2 requesting and waiting for the exclusive availability of a latch which is shareable by each

old table scanner process and which each old table scanner process obtains before commencing scanning.

25. The method of claim 22 in which the method further comprises the step of preventing the allocation of space for new or updated records in each respective one of a sequentially advancing defined set of pages during the vacate move step for the said one of the sequentially advancing defined set of pages.

26. The method of claim 22 in which an identity and status of moved records is maintained in a record identifier mapping table.

27. The method of claim 26 in which the record identifier mapping table comprises a mapping record corresponding to each of the moved records, each mapping record comprising an old record identifier for the initial position of the moved record, a new record identifier for the moved position of the moved record, and a status flag for the moved record.

28. The method of claim 27 in which the record identifier mapping table comprises a hash bucket comprising a hash on the new record identifier, and a hash table containing a reverse record identifier mapping.

29. The method of claim 22 in which the database table has an associated index and in which each fill step selects records by scanning the associated index of the database

3 table.

1 30. A method for reorganizing a database table online, the method comprising the
2 following steps:

3 a) a vacate move step to move data records from move pages in the table;

4 and

5 b) a fill move step to move data records into move pages in the table.

1 31. The method of claim 30 wherein each move step comprises the step of defining
2 temporary pointers from the original position of each moved record to the moved position
3 of the moved record.

4 32. The method of claim 31 further comprising the step of

5 c) defining selected scanners to be constrained to follow the temporary
pointers while accessing records so as to maintain data ordering for the selected scanners
wherein the move steps are synchronised with query processes launching scanners so as
to maintain the temporary pointers for use by said scanners launched by query processes.

1 33. The method of claim 32 in which the database table comprises overflow pointer
2 records and the original position of a moved record from which a temporary pointer
3 points.

1 34. A computer readable medium containing program instructions for reorganizing a

2 database table online, the program instructions comprising the steps of:

- 3 a) moving a subset of records within the database table;
- 4 b) flagging each moved record in as a reorganization record;
- 5 c) creating a reorganization pointer record for each moved record in at the
- 6 initial location of the moved record, the reorganization pointer record pointing to the new
- 7 location of the moved record; and
- 8 d) establishing scanner process constraints based on whether a scanner
- 9 process is commenced prior to or after the moving.

1 35. The computer readable medium of claim 34 wherein step d) further comprises:

2 d1) identifying table and index scanner processes commencing prior to the

3 completion of the moving as old table scanners wherein old table scanners are

4 constrained to access records in the database table by following reorganization pointers to

5 access reorganization records and wherein old table scanners are constrained to ignore

6 reorganization records when accessing records sequentially in a scan of the database

7 table; and

8 d2) identifying table and index scanner processes commencing after the

9 completion of the moving as new table scanners, wherein new table scanners are

10 constrained to ignore reorganization pointers in accessing records in the database table.

1 36. The computer readable medium of claim 35 further comprising the steps of:

2 e) deleting the reorganization pointer records and removing the

3 reorganization record flag from the records flagged after each query associated with an

old table scanner is complete; and

f) repeating steps a-d after the completion of each query associated with a new scanner until the records in the database table meet a reorganization criteria.

37. The computer readable medium of claim 36 in which step f) alternates between moving records from a one of successively defined portions of the table and moving records in a predefined order into a one of the successively defined vacated portions of the table.

38. A computer readable medium containing program instructions for reorganizing a database table online, the program instructions comprising the steps of:

a) a vacate move step;

b) a vacate clean up step;

c) a fill move step; and

d) a fill clean up step wherein each clean up step is synchronised to commence at the completion of a move step and to commence only when all queries launching scanner processes commenced before the completion of a move step have completed and each move step is synchronised to commence at the completion of a clean up step and to commence only when all queries launching scanner processes after the completion of a previous move step have completed.

39. The computer readable medium of claim 38 wherein step a) comprises:

i) moving each data record in a defined set of pages to a new location in a

3 page in the database table outside the defined set of pages;

4 ii) flagging each moved non-overflow data record as a reorganization record;

5 iii) creating a reorganization pointer record for each moved non-overflow data
6 record at the initial location of the moved record, the reorganization pointer record
7 pointing to the new location of the moved record;

8 iv) for each overflow pointer record located in the defined set of pages and
9 associated with a moved overflow data record, creating a reorganization pointer record
10 for the moved overflow data record at the location of the said reorganization pointer
11 record and flagging the moved overflow record as a reorganization record; and

12 v) flagging each overflow pointer record pointing to an overflow record
13 outside the defined set of pages to be a reorganization pointer record, and flagging each
14 of said overflow records as a reorganization record.

1 40. The computer readable medium of claim 39 wherein step b) comprises:

2 i) deleting reorganization pointers in the defined set of pages; and

3 ii) removing the reorganization record flag from each reorganization record
4 in the database table.

1 41. The computer readable medium of claim 40 wherein step c) comprises:

2 i) moving selected records to the defined set of pages such that the relative
3 ordering of the selected records matches the predefined ordering constraint;

4 ii) flagging each moved record as a reorganization record;

5 iii) creating a reorganization pointer record for each moved non-overflow

6 record at the initial location of the moved record, the reorganization pointer record
7 pointing to the new location of the moved record; and

8 iv) creating a reorganization pointer record for each moved overflow record at
9 the location of the overflow pointer record for the moved record, the reorganization
10 pointer record pointing to the new location of the moved record.

1 42. The computer readable medium of claim 41 wherein step d) comprises:

2 i) deleting reorganization pointers in the defined set of pages; and

3 ii) removing the reorganization record flag from each reorganization record
4 in the database table.

1 43. The computer readable medium of claim 42, wherein table scanner processes
2 commencing during each move step are identified as old table scanners, wherein old table
3 scanners are constrained to access records in the database table by following
4 reorganization pointers to access reorganization records and are further constrained to
5 access reorganization records sequentially by ignoring reorganization records and table
6 scanner processes commencing after the completion of a move step are identified as new
7 table scanners, wherein new table scanners ignore reorganization pointers in accessing
8 records in the database table.

1 44. A computer readable medium for reorganizing a database table online, the method
2 comprising the following steps:

- 3 a) a vacate move step to move data records from move pages in the table;
4 and
5 b) a fill move step to move data records into move pages in the table.

1 45. The computer readable medium of claim 44 wherein each move step comprises
2 the step of defining temporary pointers from the original position of each moved record
3 to the moved position of the moved record.

1 46. The computer readable medium of claim 45 further comprising the step of
2 c) defining selected scanners to be constrained to follow the temporary
3 pointers while accessing records so as to maintain data ordering for the selected scanners
4 wherein the move steps are synchronised with query processes launching scanners so as
5 to maintain the temporary pointers for use by said scanners launched by query processes.

1 47. A computer system for reorganizing a database table online, the computer system
2 comprising

- 3 means for moving a subset of records within the database table;
4 means for flagging each record moved by the means for moving a subset of
5 records as a reorganization record;
6 means for creating a reorganization pointer record for each record moved by the
7 means for moving a subset of records at the initial location of the moved record, the

8 reorganization pointer record pointing to the new location of the moved record; and
9 means for establishing scanner process constraints based on whether a scanner
10 process is commenced prior to or after moving the subset of records.

1 48. The computer system of claim 47 wherein the means for establishing scanner
2 process constraints further comprises:

3 means for identifying table and index scanner processes commencing prior to
4 moving the subset of records as old table scanners wherein old table scanners are
5 constrained to access records in the database table by following reorganization pointers to
6 access reorganization records and wherein old table scanners are constrained to ignore
7 reorganization records when accessing records sequentially in a scan of the database
8 table; and

9 means for identifying table and index scanner processes commencing after the
10 completion of moving the subset of records as new table scanners, wherein new table
11 scanners are constrained to ignore reorganization pointers in accessing records in the
12 database table.

1 49. The computer system of claim 48 further comprising:

2 means for deleting the reorganization pointer records and removing the
3 reorganization record flag from the records by the means for flagging after each query
4 associated with an old table scanner is complete; and

5 means for moving another subset of records after the completion of each query
6 associated with a new scanner until the records in the database table meet a

7 reorganization criteria.

1 50. The computer system of claim 49 in which the means for moving another subset
2 of records alternates between moving records from a one of successively defined portions
3 of the table and moving records in a predefined order into a one of the successively
4 defined vacated portions of the table.

1 51. A computer system for reorganizing a database table comprising:
2 means for performing a vacate move step;
3 means for performing a vacate clean up step;
4 means for performing a fill move step; and
5 means for performing a fill clean up step wherein each clean up step is
6 synchronised to commence at the completion of a move step and to commence only when
7 all queries launching scanner processes commenced before the completion of a move step
8 have completed and each move step is synchronised to commence at the completion of a
9 clean up step and to commence only when all queries launching scanner processes after
10 the completion of a previous move step have completed.

1 52. The computer system of claim 51 wherein the means for performing the vacate
2 move step comprises:

3 means for moving each data record in a defined set of pages to a new
4 location in a page in the database table outside the defined set of pages;

5 means for flagging each moved non-overflow data record as a reorganization

6 record;

7 means for creating a reorganization pointer record for each moved non-overflow
8 data record at the initial location of the moved record, the reorganization pointer record
9 pointing to the new location of the moved record;

10 means for each overflow pointer record located in the defined set of pages and
11 associated with a moved overflow data record, creating a reorganization pointer record
12 for the moved overflow data record at the location of the said reorganization pointer
13 record and flagging the moved overflow record as a reorganization record; and

14 means for flagging each overflow pointer record pointing to an overflow record
15 outside the defined set of pages to be a reorganization pointer record, and flagging each
16 of said overflow records as a reorganization record.

1 53. The computer system of claim 52 wherein the means for performing the vacate
2 clean up step comprises:

3 means for deleting reorganization pointers in the defined set of pages; and

4 means for removing the reorganization record flag from each reorganization
5 record in the database table.

1 54. The computer system of claim 53 wherein the means for performing the fill move
2 step comprises:

3 means for moving selected records to the defined set of pages such that the
4 relative ordering of the selected records matches the predefined ordering constraint;

5 means for flagging each moved record as a reorganization record;

6 means for creating a reorganization pointer record for each moved non-overflow
7 record at the initial location of the moved record, the reorganization pointer record
8 pointing to the new location of the moved record; and

9 means creating a reorganization pointer record for each moved overflow record at
10 the location of the overflow pointer record for the moved record, the reorganization
11 pointer record pointing to the new location of the moved record.

1 55. The computer system of claim 54 wherein the means for performing the fill clean
2 up step comprises:

3 means for deleting reorganization pointers in the defined set of pages; and
4 means for removing the reorganization record flag from each reorganization
5 record in the database table.

1 56. The computer system of claim 55 further comprising:

2 means for identifying table scanner processes commencing during the
3 performance of each move step as old table scanners, wherein old table scanners are
4 constrained to access records in the database table by following reorganization pointers to
5 access reorganization records and are further constrained to access reorganization records
6 sequentially by ignoring reorganization records; and

7 means for identifying table scanner processes commencing after the completion of
8 the performance of a move step as new table scanners, wherein new table scanners ignore
9 reorganization pointers in accessing records in the database table.

1 57. A computer system for reorganizing a database table online, the computer system
2 comprising:

3 a) means for carrying out a vacate move step to move data records from
4 move pages in the table; and

5 b) means for carrying out a fill move step to move data records into the move
6 pages in the table.

1 58. The computer system of claim 57 further comprising:

2 means for defining temporary pointers from an original position of each moved
3 record to the moved position of the moved record;

4 means for defining selected scanners to be constrained to follow the temporary
5 pointers while accessing records so as to maintain data ordering for the selected scanners;
6 and

7 means for synchronising the move steps with query processes launching scanners
8 so as to maintain the temporary pointers for use by said scanners launched by query
9 processes.

1 59. The computer system of claim 58 in which the database table comprises overflow
2 pointer records and the computer system further comprises means for setting the original
3 position of a moved record to be the position of an overflow pointer record, from which a
4 temporary pointer points, where the moved data record is originally pointed to by said
5 overflow pointer record.